Supporting Mobile Multimodal Interaction with a Rule-Based Framework

Andreas Möller [1], **Stefan Diewald** [1], **Luis Roalter** [1], **Matthias Kranz** [2]
[1] Technische Universität München, Arcisstraße 21, 80333 Munich, Germany
[2] Universität Passau, Innstraße 43, 94032 Passau, Germany
andreas.moeller@tum.de, stefan.diewald@tum.de, roalter@tum.de, matthias.kranz@uni-passau.de

arXiv:1406.3225v1 [cs.HC] 12 Jun 2014

ABSTRACT
Multimodality can make (especially mobile) device interaction more efficient. Sensors and communication capabilities of modern smartphones and tablets lay the technical basis for its implementation. Still, mobile platforms do not make multimodal interaction support trivial. Building multimodal applications requires various APIs with different paradigms, high-level interpretation of contextual data, and a method for fusing individual inputs and outputs. To reduce this effort, we created a framework that simplifies and accelerates the creation of multimodal applications for prototyping and research. It provides an abstraction of information representations in different modalities, unifies access to implicit and explicit information, and wires together the logic behind context-sensitive modality switches. In the paper, we present the structure and features of our framework, and validate it by four implemented demonstrations of different complexity.

Author Keywords
Multimodality; Multimodal Interaction; Mobile Interaction; Context Awareness; Framework; Android

INTRODUCTION AND MOTIVATION

An emerged concept in the context of (mobile) interaction is *multimodality* [23], i.e., using more than one modality (simultaneously or sequentially) for input and output. This covers not only tactile, voice or gesture-based interaction, but also more implicit information, which is gathered from, e.g., the multitude of sensors that are nowadays available in mobile devices. Particularly in mobile settings, individual unimodal interaction modes can suffer from limitations (e.g., the *small screen* problem), so that multimodality may have a synergetic effect and may contribute to more efficiency [23]. Furthermore, human perception and communication uses all senses, so that multimodal interaction is considered more natural [12] and thus more intuitive [30].

Multimodality support is beneficial in various individual uses cases. One example is in-car interaction [5–7], where auditory in- and output and gestures [24] support hands-free interaction. One-handed interaction is e.g. supported by distance-based input methods like DistScroll [14].

Often, the requirements for choosing a certain modality depend on the context, e.g., time, location, social setting, or security demands [26]. As an example for context-driven modality settings, Alice may not want to be notified on emails on weekdays between 10 PM and 6 AM; and Bob may mute his private phone at work, and always switch to hands-free mode in the car. For each of those exemplified situations, a certain modality change (either on input or output side) occurs. While these contextual modality switches mostly have to be performed manually (e.g., muting the phone before the meeting begins), they often occur in recurring situations, which are often based on rather simple rules. Most of the context types that play a role for modality decisions could even be inferred by the mobile device itself.

However, today's solutions address rather particular problems (take e.g. the iPhone's *Do Not Disturb* mode that pauses notifications in a definable period of time), which are either not sufficiently flexible, or have to be bought as additional software for each particular problem. Making use of multimodality is still cumbersome for mobile software developers, although the necessary functionality is present in current devices. Yet, accessing and fusing contextual and sensor data in a meaningful way to create context-sensitive applications that feel "smart" is still a considerable effort. End users do not have the possibility at all to create context-based behavior fitted to their personal needs.

Besides context-based switching of output channels, multimodality also allows new input methods and paradigms. Novel interaction modalities have shown superior performance in e.g. task completion time than traditional methods [18] or enable novel scenarios in different domains such as indoor navigation [20], health and fitness [21], or education and edutainment [11, 22]. However, the adoption of multimodality in the real world is still cautious, both from user and programmer side.

In this work, we present the *M3I* (*M*obile *M*ulti*M*odal *I*nteraction) Framework [19], to our knowledge the first framework to support rich, context-driven multimodal interaction in mobile applications in a fast, holistic and simple manner. It allows researchers and application developers to integrate various modalities and to set up the decision logic for switching between them based on simple rules, all with a handful lines of code. Explicit interaction is supported as well as implicit, context-driven behavior. No special setup and development tools are required, and the framework is fully extensible thanks to its modular structure.

The paper is organized as follows. We conducted a focus group and expert interviews to deeper investigate current limitations and wishes, both from a user and a software developer point of view. We briefly discriminate our framework

from related research in multimodality and context awareness toolkits. Building upon the requirements analysis, we then present the *M3I* framework and its components, illustrated by code examples. We validate the claim that the *M3I* framework drastically simplifies the support of multimodal interaction in mobile applications as well as the integration of context information. We show this at the example of several different use cases that demonstrate the broad applicability of the framework. Finally, we discuss possible enhancements of the system for future work.

REQUIREMENTS ANALYSIS

As an initial step, we conducted a focus group and expert interviews to identify requirements both from a user's and a developer's perspective.

User's Perspective: Focus Group

The focus group was conducted with six participants (5 males, 1 female) between 24 and 30 years (average age: 27). Four participants were research assistants; two were students. Four of them owned an Android smartphone and two had an iPhone. All participants rated their smartphone expertise as high or very high. In a guided discussion, we investigated current usage and adoption of multimodality, and elicited unsatisfied needs and expectations. The focus group took about one hour and was audio-recorded. In the following, we summarize the most important results and design implications.

Current usage of input and output modalities

While the touchscreen is the prevailing interaction method, the usage of further modalities broadly varied between participants. One participant used vibration and notification lights as preferably used output; another participant did not use vibration at all, but heavily relied on speech input (although only in private space). A third participant neither used sound nor vibration, but relied on screen notifications, keeping the phone next to him on the desk. This indicates that output modalities should account for diverse preferences and defaults. Further, the focus group revealed that modalities are rarely changed, but that participants maintain default settings which they alter only in certain situations. One subject enabled sound only when expecting an important call, while another had the phone in ringing mode as default, but muted the phone in silent environments (e.g., the library). This shows that defaults are interpreted differently.

It also turned out that when speaking of modalities, participants had mostly output in mind. In terms of input methods (other than the touchscreen), participants stated to use device buttons to control the music application or to decline calls, and voice input to set a timer or perform a search query. Hence, input modalities are rather task-specific; therefore a frame-work should support the implementation of novel modalities for individual purposes.

Identified Problems in Status Quo

Subjects stated that alternative input methods (i.e., other than standard touch interaction) either do not work reliably enough (speech) or are not implemented consistently (e.g., tap pattern gestures or gaze-based interaction, as offered by some Samsung phones). Participants desire that interaction methods should be more consistent and work in a similar manner in all applications. One participant stated that he often forgot to revoke a modality change, which could result in unwanted situations (e.g., the phone rings in a silent environment). This fear was confirmed by almost all subjects and pointed out as reason for the presently rare changes of modalities. They would welcome a system that automatically selects modalities, but would like to retain control and be able to override the system behavior. Participants also missed a solution similar to "profiles" (as available on earlier Nokia phones), so that a rule-based modality switching approach seemed attractive, especially using context information as a basis for decisions. Participants supposed that most situations could possibly be covered by a small number of rules. Users could then define their own modality settings, which are activated by conditions (such as location, time, etc.) and revoked when the conditions do not apply any more. Participants supposed that rules would be often initially set and later rarely modified. They pointed out that users always need the possibility to override settings when a rule becomes active. In order to prevent forgetting, another idea would be a time limitation for manual settings (e.g., one hour or until the meeting is finished).

Developer's Perspective: Expert Interviews

For the programmer's perspective, we interviewed three software developers involved in mobile application development and asked how satisfied they were with the current tool support for creating multimodal applications. We also were interested in their expectations and wishes for future tools regarding their programming needs.

The following issues were mentioned (summarized and aggregated): Implementing contextual behavior requires the use of different APIs (e.g., Location API, Sensor API, etc.), entailing redundancies and re-implementations of similar pieces of code, as well as heterogeneous ways of accessing data. While for example sensor or location updates are listener-based (push principle), other data must be checked manually (pull principle). To realize context-sensitive behavior based on both push- and pull-based data, interfaces/wrappers need to be created, which adds significant overhead. Developers would appreciate a unified structure for all types of context information, as well as encapsulation of frequently used functions, hiding complexity. In terms of input, novel interaction methods currently have to be designed and implemented from scratch. Especially for rapid prototyping, building blocks would speed up the creation of functional prototypes. In terms of output, it is effortful to include multiple modalities in an application, as each additional output modality must be implemented separately. If developers could in their implementation abstract from the information to be communicated and the method/channel over which it is transported, they were more likely to implement multiple modalities, contributing to more usable, natural, intuitive, and efficient applications.

Summary of Requirements

Based on focus group and expert interviews, we can summarize the following requirements:

- A holistic system for mobile multimodal interaction should cover both multimodal input (in terms of natural interaction) and output (in terms of context-based modality switches).
- Defining multimodal behavior based on rules is closest to the human understanding of automated switching and might thus be an adequate underlying model.
- The system must be flexible enough for the heterogeneous preferences in terms of situational modality settings and defaults.
- The underlying programming framework should save time, reduce code redundancy, ease the access of contextual information, and abstract from information representations (so that a unit of information can be easily communicated by different modalities).

DISCRIMINATION FROM RELATED WORK

For rapid development of context-sensitive applications in the research context, several toolkits and frameworks have been presented, starting with the *Context Toolkit* by Dey et al. [4]. The *Subtle* toolkit by Fogarty and Hudson [10] supports the development of applications that use sensor-based statistical models. Thereby, contextual information can be used to adapt certain settings automatically. With our *M3I* framework, we focus on *rule-based* wiring of functionality, as often the impacting factors for a desired action are quite clear. There exist alternative approaches (e.g., support vector machines, Bayesian or neural networks), but machine learning often makes it hard to understand *why* a certain action has been performed. Some toolkits address particular use cases, e.g. physical mobile interaction [27] or proxemic interaction [17]. Du and Wang [8] developed a model and implementation framework to simulate context events with Symbian phones. Schuster et al. used JCOP (Context-Oriented Programming) extensions to the Java language that support context-dependent execution for Android programming [29], which, however, complicates a quick integration in any existing Android project and working environment.

Preceding toolkits and frameworks for multimodality that have been presented in research are often focused and confined to specialized use cases, e.g., speech and gesture interaction with large screens [13], or multimodal interaction on a PC [2, 9]. The generalizability of such approaches is thus limited. The evaluation of multimodal behavior is conducted by e.g. finite automata [3] or state machines [1]. The latter approach enhances the Java Swing toolkit to facilitate novel input methods (e.g., multi-handed or pressure-sensitive input). Ravindranath et al. [25] showed a server-based task distribution and coordination approach, which however entails that the systems using it are not fully autonomous. Manca et al. [16] present a multimodal web toolkit, however focusing on browser-based applications. These have limited capabilities compared to native applications, e.g., with relation to hardware and sensor access on mobile devices.

For the choice of optimal modalities (using e.g. the visual, haptic, and auditory channel), often also the users context, such as time, location, or the social setting, plays an important role. Context determination is supported by different approaches, such as toolkits [4], middleware [15] or context-oriented programming extensions on code level [29]. They facilitate e.g. tasking applications like Tasker[1] that utilize context to change output modalities (for example, mute ringtones in a silent environment). While this overview does not claim to be complete, existing toolkits and frameworks seem not to match all of the requirements that emerged from the focus group (see "Requirements Analysis" section). To the best of our knowledge, the *M3I* framework we present in this paper is currently the only approach that fulfills all of the following points:

- Holistic support of modeling, prototyping and implementation of multimodal behavior (often, only individual stages, such as the prototyping stage, are supported)
- Support of both multimodal input and output (other approaches focus either on the interpretation and fusion of input events, or on output modality switching in form of tasking applications)
- Support of context-driven multimodality (existing approaches are either pure context toolkits, or focus on multimodal interaction without integrating context, at least not using context for both in- and output)
- Applicability in standard development environments (code enhancements like JCOP [29] require adaptations such as a special compiler)
- Applicability with state-of-the-art mobile operating systems like Android (earlier approaches are dedicated to desktop computing or legacy platforms with limited support for context acquisition)
- Autonomous operation, without server-based evaluation or control through the cloud as e.g. in Code in the Air [25]
- Compliance with the human mental model of multimodal behavior (systems that use automata, state machines, etc. might have worse intelligibility than a rule-based system)

M3I: MOBILE MULTIMODAL INTERACTION FRAMEWORK

Our *M3I* framework is implemented as Android library (using the Java language). Thereby, no special configuration or tools are needed for using it; the library just needs to be referenced from an Android project to use its features. Figure 1 shows, on a conceptual level, the structure of the framework and its basic components, which will be detailed in the following.

Context Factors and Context Groups

Context factors are fine-grained pieces of context information that can be used to influence the behavior of a context-sensitive application. Examples for context factors are the latitude and longitude of the device's position, the charging state of the battery, or the time of day, but also complex information such as the user's current activity. In the framework, a

[1] `https://play.google.com/store/apps/details?id=net.dinglisch.android.taskerm`, accessed June 4 , 2014

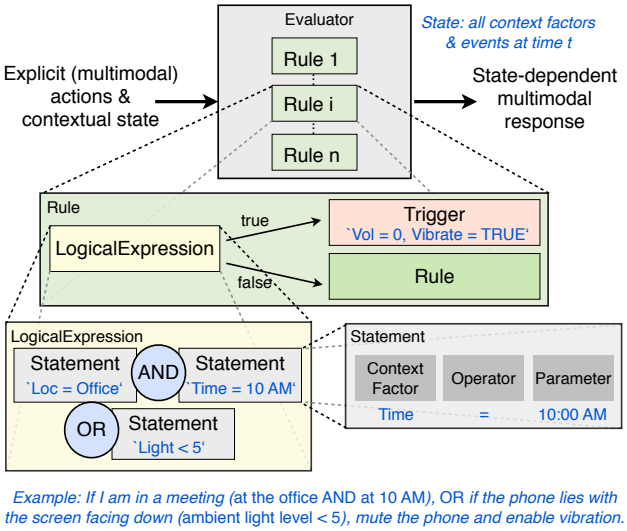

Figure 1. General structure of the *M3I* framework: The core is the evaluator, which evaluates rules that define the system's behavior on explicit user actions and/or implicit context factors and events. Rules can initiate triggers to directly react on the context and to, e.g., switch modalities, and recursive rules allow implementing complex decision-making.

`ContextFactor` is defined by specifying a `ContextGroup` it belongs to, and the respective context method that provides the `ContextFactor`'s value. Since Java does not support function references, constants specify the method ID. They are evaluated by an `execute()` method to call the context function appropriately at runtime.

For example, a context factor of type `Boolean` indicating the device's charging level is defined as follows:

```
FloatContextFactor battLevel =
  new FloatContextFactor(
    new BatteryContext(this),
    BatteryContext.FLOAT_GET_BATTERY_LEVEL);
```

Methods for retrieving context information are organized in groups such as `OrientationContext` or `LightContext`, which on their part contain information on the device's orientation or the ambient light level. Context groups can be seen as collections that organize and provide access to the platform's own methods and routines in a way that eases building up statements and rules (see the following sections).

The framework drastically economizes and simplifies those method calls for the programmer, as it saves all overhead for initializing system services, creating event listeners etc., and makes all functionality accessible by homogenous interfaces. One particular advantage of the framework is unified handling of synchronous and asynchronous information. Values available anytime (e.g., the current weekday) can be retrieved with a simple method call, but dynamic events are normally handled with listeners to receive them asynchronously when they occur (e.g., location changes, sensor events, or touch interactions by the user). The framework makes transparent to the developer whether a `ContextFactor` is based on a synchronous or an asynchronous call. It internally creates listeners if required, handles their updates automatically and stores the most recent values. Thus, whenever the contextual state should be determined, a consistent state is available.

Abstraction of Information Representation by Triggers

Triggers define what should happen at a defined contextual state. They realize modality switches, i.e., they abstract how information shall be represented. For example, a notification can be received as sound or as vibration signal. A row of modalities are directly supported, e.g. visual (UI changes), haptic (vibration patterns), or auditory (sound playback) responses. Triggers are are not limited to predefined modalities or actions: a `MethodTrigger` calls an arbitrary method to implement custom functionality. With the `NullTrigger` the omission of an action, e.g. in an *else* branch, can be modeled. For an example of how triggers are defined, see Demonstration 1 in the *Validation* section.

Rule-Based Wiring

The wiring of input and contextual events and appropriate triggers is realized based on rules. A `Rule` consists of a `LogicalExpression` which is evaluated to be either true or false, and of actions that define what should happen in either case. `Rules` can either call a `Trigger` (as described above), or recursively another `Rule`, which allows the creation of a more complex, nested decision logic.

Logical expressions

The simplest form of a logical expression is a `Statement`, which consists of a `ContextFactor` and an `Operator`. Operators allow value checks such as numeric comparisons, within-range-tests, or regular expressions. Take for example the following statement which checks whether the battery level is above 50% (it uses the `ContextFactor battLevel` defined in a previous example):

```
Statement isAboveHalfCharged = new Statement(
  battLevel, FloatOperator.greaterThan(50f));
```

Using logical operators realized by `UnaryExpression` and `BinaryExpression` classes, complex terms can be created Although AND, OR, and NOT are sufficient to construct any logical expression according to Boolean algebra laws, XOR, NAND, NOR and XNOR are supported for convenience reasons to implement a decision logic of arbitrary complexity.

The following expression describes the state in which the device is either charged more than 50% *or* connected to a power plug (in the example, `isAboveHalfCharged` and `isPluggedIn` are previously defined statements):

```
BinaryExpression exp = new BinaryExpression(
  BinaryExpression.EXPRESSION_OR,
  isAboveHalfCharged, isPluggedIn);
```

When triggers are defined and the contextual states in which they should be applied have been described by logical expressions, a rule can be created.

Rule Evaluation

The set of active rules is evaluated in the framework's `StateMachine`. The developer can define the update interval

in which the evaluator evaluates the rules and executes the defined triggers. Thereby, the framework supports time-critical, fast-reacting systems as well as battery-conserving systems where an update e.g. once per minute is sufficient. Once a few rules have been defined, the following lines of code activate the framework in any Android application (using a 1000 milliseconds update interval in this example):

```
StateMachine m = new StateMachine(1000);
m.addRule(r_1);
...
m.addRule(r_n);
m.start();
```

Extensibility

With a `CustomContext` group, results of any method or callback can be fed into the framework's decision logic when a context factor is not built in directly. Moreover, by clearly defined interfaces, the framework is easy to extend both on input and output side. The following simple steps are e.g. required to add a new context group. (1) Create a new class implementing the interface `IContextGroup`, and implement all context methods the group should provide, returning a value of arbitrary data type. (2) Define a list of parameters for each function that serves as argument for the `execute()` function of the group. In a similar manner, new triggers must implement the interface `ITrigger`, which contains a `trigger()` method that performs the desired action.

Framework Summary

The framework currently integrates more than 50 context factors regarding, e.g., location, ambient noise and light level, device orientation, battery information, proximity information (through NFC, Bluetooth, or Geofence entering/leaving), availability of 3G and WLAN connections, or date and time. Basic activity recognition and classification routines abstracting from pure sensor readings are already integrated, e.g. pose classification (in pocket or carried in hand), usage indicators, mode of transportation, vision-based detection (face) etc.

Besides that, explicit interactions, such as physical button presses or touch interactions can be intercepted and combined with implicit contextual information.

On output side, triggers allow controlling a range of modalities and thereby abstracting how information should be represented, based on fine-grained definitions. Examples include visual output (on-screen or via LEDs), sound, brightness setting, or vibration. Further actions include behavioral triggers, e.g. changing device settings (sync rules, connectivity, screen lock, ...) or custom functions.

VALIDATION

To validate our claim that the framework drastically speeds up the development of context-based multimodal applications, we present three very different applications demonstrating the wide spectrum of potential use cases for our framework.

Demonstration 1: Flip to Mute

This basic example is a tangible "mute by flip over" application for quickly enabling silent mode, e.g. in a meeting. Instead of manually turning down the volume, the user simply needs to flip over the phone and place it on the table with the display facing down (see Image 1 in Figure 2). A dozen years ago, such context-based telephony applications were prototyped with external sensor modules [28]. In our example, the trigger is realized using the built-in ambient light sensor mounted on the front panel of the phone. With our framework, this is the code to set up this functionality:

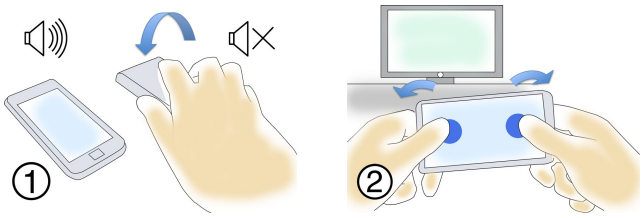

Figure 2. Multimodal interaction examples that can be realized with few lines of code using our framework: (1) A turn-to-mute service, switching the phone to silent mode when it is placed upside down (detected by the ambient light sensor). (2) A multimodal game controller using the accelerometer to steer a video game running on a second screen.

```
// define statement
FloatContextFactor light =
   new FloatContextFactor(lc,
   LightContext.FLOAT_GET_LIGHT_LEVEL);
Statement isUpsideDown = new Statement(light,
     FloatOperator.smallerThan(5.0f));
// define triggers
AudioTrigger mute = new AudioTrigger(this);
mute.setAction(AudioTrigger.RINGER_VIBRATE);
AudioTrigger ring = new AudioTrigger(this);
ring.setAction(AudioTrigger.RINGER_NORMAL);
// create rule
Rule r = new Rule(isUpsideDown, mute, ring);
```

The example does not take into account that the light level can be 0 also at night or when the phone is in a sleeve (it could easily be refined by adding the time of day and the pose as additional context factors), but it demonstrates how few lines of code suffice to realize a multimodal interaction.

Demonstration 2: Multimodal Game Controller

The second example shows how to use a mobile phone as game controller, combining the modalities *touch* and *gestures*. In a racing game running on a second screen, the vehicle can be steered by tilting the phone (see Image 2 in Figure 2). The angle of the steering wheel is derived from the pose detection context factor, which can directly retrieve the tilt angle of the device. Touching additional action buttons in the interface on the mobile screen can modify the driving behavior, which demonstrates the framework's capability to seamlessly integrate explicit (button presses) and implicit events (updates from the orientation sensor listener). The communication to the game running on an Ubuntu desktop PC is realized using ROSjava[2], a publish/subscribe architecture based on the Robot Operating System[3], a popular architecture for distributed applications in research.

[2] **`http://code.google.com/p/rosjava/`**, accessed June 4, 2014

[3] **`http://www.ros.org`**, accessed June 4, 2014

Input Method	Description
	Raise&Call *(to launch Phone app)* The user raises the phone to his ear and speaks the name of the person to call. *Mimicked movement: answering a call* *Involved Modalities: Motion Gesture, Speech*
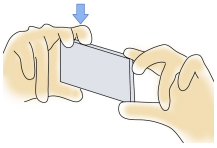	**Press&Shoot** *(to launch Camera app)* The user brings the phone to her front in an upright position and presses the volume button. *Mimicked movement: taking a photo* *Involved Modalities: Motion Gesture, Button Press*
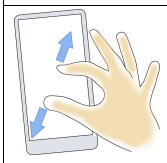	**Pinch@Home** *(to launch Maps app)* With her thumb and index finger, the user performs a multi-touch pinch gesture on the home screen. *Mimicked movement: enlarging an area in a map* *Involved Modalities: Screen Gesture*

Figure 3. Individual intuitive input methods defined with the *M3I* framework to start various applications, mimicking typical gestures and movements when using these apps.

Demonstration 3: Mimikry Input Methods

As multimodal input is (as argued earlier), in contrast to output, task-specific, we implemented three unique methods (*Raise&Call*, *Press&Shoot*, and *Pinch@Home*) to launch different applications (Phone, Camera, Maps). As our goal was high intuitiveness, we chose metaphors that support the human mental model of performing typical movements with these apps. Each method includes different modalities, like motion gestures (performing a characteristic motion with the device), screen gestures (performing a multi-touch gesture on the screen), or explicit voice or button input (see first three rows of Table 1). These three input methods were implemented with *M3I* in straightforward manner by a few simple rules (e.g., involving the accelerometer pose as implicit sensor reading and the button press as explicit action for the Press&Shoot method). Built-in abstractions (predefined device poses like "upright", "lying on the table", "display up/down") additionally accelerated the development of the *Raise&Call* and *Press&Shoot* methods. With the rule creation interface shown in Figure 2, end users can easily define similar interaction methods on their own.

Demonstration 4: End User Toolkit

In a final example, we do not focus on a particular scenario, but show an all-purpose application that provides access to the entire functionality of the framework. As motivated in the introduction, users often have very few and concrete factors that determine a desired modality. However, it is cumbersome to toggle the respective settings each time per hand. With help of our framework it was little effort to assemble an application allowing end users to automate such actions (similar to, e.g., Llama[4], or Locale[5] for Android, which are however not

[4] `https://play.google.com/store/apps/details?id=com.kebab.Llama`, accessed June 10, 2014

[5] `http://www.twofortyfouram.com`, accessed June 10, 2014

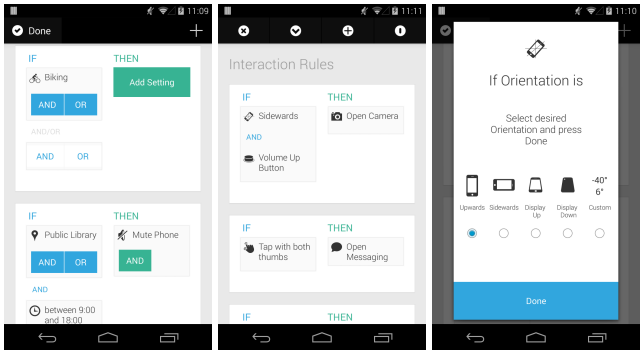

Figure 4. With a graphical user interface (GUI), end users can make use of the entire functionality of the *M3I* framework and set up their individual context-defined multimodal behavior. Left: Definitions of new multimodal input methods. Middle: Editing rules for context-dependent output modality switching. Right: Recording orientation (template-based or custom) as part of a multimodal input method.

focusing on multimodal interaction). Using a GUI, users can specify their rules conveniently to, e.g., define gestures that trigger certain interactions, contexts in which certain modalities should be enabled, and more. Based on the GUI, the application creates all required logical expressions, triggers and rules, and runs the evaluator as background service. Unlike existing applications which confine to a predefined set of simple actions and rules, our framework provides a complex decision logic, high-level contextual abstractions with activity recognition, and it is easily extensible: Recompiling the application with an updated version of the framework directly makes newly added functions accessible.

Figure 4 shows our implementation of a rule-based interface to define system-wide multimodal behavior based on contextual conditions. This interface internally creates `LogicalExpression`s, `Trigger`s and `Rule`s, but is intuitive enough to be operated by end users. In concordance with a suggestion made in the focus group, the system goes back to the previous setting when a rule does not apply any more. This allows using any setting as default. While the left screenshot shows the definitions of some of the interaction methods presented in Demonstration 3, the middle image is an example for the creation of multimodal behavior depending on different contexts, like the mode of transportation or the location. The right screenshot shows the recording of orientation gestures as part of a rule for multimodal interaction.

CONCLUSION AND FUTURE WORK

The presented *M3I* framework gives researchers and developers a tool at hand that simplifies rapid development and prototyping of multimodal interaction in mobile applications. While still under active development, we have demonstrated its maturity by the examples in this paper.

In future work, besides adding more context factors, we will experiment with automated deduction of desired modalities. The framework could silently observe user behavior, identify patterns and suggest new rules. For example, when a user mutes the device every week at a certain location and time (probably because of a weekly meeting), it could offer to perform this automatically. Unlike classical supervised learning

where the learned model is not visible to the user, the rules to be created could be edited and reviewed to establish transparency about an application's multimodal behavior.

Acknowledgment
We thank our student Max Walker for his help with extending the *M3I* framework and with creating the graphical end user toolkit on top of the framework.

REFERENCES

1. Appert, C., and Beaudouin-Lafon, M. SwingStates: Adding state machines to Java and the Swing toolkit. *Software: Practice and Experience 38*, 11 (2008), 1149–1182.
2. Bourguet, M.-L. A toolkit for creating and testing multimodal interface designs. In *Proceedings of the ACM User Interface Software and Technology Symposium (UIST 2002)*, ACM (2002), 29–30.
3. Cutugno, F., Leano, V. A., Rinaldi, R., and Mignini, G. Multimodal framework for mobile interaction. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, ACM (2012), 197–203.
4. Dey, A. K., Abowd, G. D., and Salber, D. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction 16*, 2 (Dec. 2001), 97–166.
5. Diewald, S., Möller, A., Roalter, L., and Kranz, M. DriveAssist - A V2X-Based Driver Assistance System for Android. In *Mensch & Computer Workshopband*, H. Reiterer and O. Deussen, Eds., Oldenbourg Verlag (2012), 373–380.
6. Diewald, S., Möller, A., Roalter, L., and Kranz, M. MobiliNet: A Social Network for Optimized Mobility. In *Adjunct Proceedings of the 4th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI 2012)* (2012), 145–150.
7. Diewald, S., Möller, A., Roalter, L., Stockinger, T., and Kranz, M. Gameful design in the automotive domain: review, outlook and challenges. In *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI 2013)*, ACM (2013), 262–265.
8. Du, W., and Wang, L. Context-aware application programming for mobile devices. In *Proceedings of the 2008 C3S2E conference*, ACM (2008), 215–227.
9. Flippo, F., Krebs, A., and Marsic, I. A framework for rapid development of multimodal interfaces. In *Proceedings of the 5th International Conference on Multimodal Interfaces (ICMI 2003)*, ACM (2003), 109–116.
10. Fogarty, J., and Hudson, S. E. Toolkit support for developing and deploying sensor-based statistical models of human situations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2007)*, ACM (2007), 135–144.
11. Holleis, P., Kranz, M., Winter, A., and Schmidt, A. Playing with the real world. *Journal of Virtual Reality and Broadcasting 3*, 1 (2006), 1–12.
12. Jokinen, K., and Hurtig, T. User expectations and real experience on a multimodal interactive system. In *Proceedings of the INTERSPEECH* (2006).
13. Krahnstoever, N., Kettebekov, S., Yeasin, M., and Sharma, R. A real-time framework for natural multimodal interaction with large screen displays. In *Proceedings of the 4th IEEE International Conference on Multimodal Interfaces (ICMI)*, IEEE (2002), 349.
14. Kranz, M., Holleis, P., and Schmidt, A. DistScroll – A New One-Handed Interaction Device. *25th IEEE International Conference on Distributed Computing Systems Workshops 5* (2005), 499–505.
15. Lee, Y., Iyengar, S., Min, C., Ju, Y., Kang, S., Park, T., Lee, J., Rhee, Y., and Song, J. MobiCon: a mobile context-monitoring platform. *Communications of the ACM 55*, 3 (2012), 54–65.
16. Manca, M., Paternò, F., Santoro, C., and Spano, L. D. Generation of multi-device adaptive multimodal web applications. In *Mobile Web Information Systems*. Springer, 2013, 218–232.
17. Marquardt, N., Diaz-Marino, R., Boring, S., and Greenberg, S. The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology (UIST 2011)*, ACM (2011), 315–326.
18. Möller, A., Diewald, S., Roalter, L., and Kranz, M. MobiMed: comparing object identification techniques on smartphones. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design (NordiCHI 2012)*, ACM (2012), 31–40.
19. Möller, A., Diewald, S., Roalter, L., and Kranz, M. M3I: A Framework for Mobile Multimodal Interaction. In *Proceedings of Mensch & Computer 2014: Interaktiv unterwegs – Freiräume gestalten*, Oldenbourg Verlag (München, 2014).
20. Möller, A., Kranz, M., Huitl, R., Diewald, S., and Roalter, L. A mobile indoor navigation system interface adapted to vision-based localization. In *Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia (MUM 2012)*, ACM (2012), 4:1–4:10.
21. Möller, A., Roalter, L., Diewald, S., Scherr, J., Kranz, M., Hammerla, N., Olivier, P., and Plötz, T. Gymskill: A personal trainer for physical exercises. In *IEEE International Conference on Pervasive Computing and Communications (PerCom 2012)* (2012), 213 –220.

22. Möller, A., Thielsch, A., Dallmeier, B., Roalter, L., Diewald, S., Hendrich, A., Meyer, B. E., and Kranz, M. MobiDics – improving university education with a mobile didactics toolbox. In *Ninth International Conference on Pervasive Computing (Pervasive 2011), Video Proceedings* (2011).

23. Oviatt, S. Ten myths of multimodal interaction. *Commun. ACM 42*, 11 (Nov. 1999), 74–81.

24. Pfleging, B., Kern, D., Döring, T., and Schmidt, A. Reducing non-primary task distraction in cars through multi-modal interaction. *it - Information Technology 54*, 4 (2012), 179–187.

25. Ravindranath, L., Thiagarajan, A., Balakrishnan, H., and Madden, S. Code in the air: simplifying sensing and coordination tasks on smartphones. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, ACM (2012), 4.

26. Riedl, P., Koller, P., Mayrhofer, R., Möller, A., Koelle, M., and Kranz, M. Visualizations and switching mechanisms for security zones. In *Proceedings of International Conference on Advances in Mobile Computing & Multimedia (MoMM 2013)*, ACM (2013), 278:278–278:281.

27. Rukzio, E., Broll, G., and Wetzstein, S. The Physical Mobile Interaction Framework (PMIF). Tech. Rep. LMU-MI-2008-2, Ludwig-Maximilians University Munich, 2008.

28. Schmidt, A., Aidoo, K. A., Takaluoma, A., Tuomela, U., Laerhoven, K. V., and Velde, W. V. d. Advanced Interaction in Context. In *Proceedings of the 1st Intl. Symposium on Handheld and Ubiquitous Computing (HUC 1999)*, Springer (1999), 89–101.

29. Schuster, C., Appeltauer, M., and Hirschfeld, R. Context-oriented Programming for Mobile Devices: JCOP on Android. In *Proceedings of the 3rd International Workshop on Context-Oriented Programming (COP 2011)*, ACM (2011), 5:1–5:5.

30. Watanuki, K., Sakamoto, K., and Togawa, F. Analysis of multimodal interaction data in human communication. In *Proceedings of the International Conference on Spoken Language Processing (ICSLP 1994)*, vol. 94 (1994), 899–902.